Virginia Polytechnic Institute and State University

Electrical Engineering
BLACKSBURG, VIRGINIA 24061

TOWARD THE ASSESSMENT OF THE SUSCEPTIBILITY OF A DIGITAL SYSTEM TO LIGHTNING UPSET

by

Joseph G. Tront
Werner Graf
Clyde F. Martin
Gerald M. Masson
John M. Myers
James J. Whalen

December 1981

Final Report
Grant NAG 1-132, "Lighting Upset Methodology Research"

Prepared for

Flight Electronics Division
NASA Langley Research Center
Hampton, Virginia

CONTENTS

I. Aspects of Lightning Upset

A. Electromagnetic environment

1. Lightning stress

2. Shielding

B. System properties and effects

1. Stress characterization; parameters TV and TR

2. System characterization

3. Upset detection

4. Lists of positive and negative design features

II. First-cut Theory of Comparing Candidate Designs

A. Tests of comparative susceptibility

B. Analysis and simulation of comparative susceptibility

III. An Approach to Lightning-Induced Transient Fault Effects

IV. Tasks for Future Work

Toward the Assessment of the Susceptibility of a Digital System to Lightning Upset

We report on accomplishments and directions for further research aimed at developing methods to assess a candidate design of an avionic computer, especially a fault-tolerant computer, with respect to susceptibility to lightning upset. This report is presented in three main sections. Section I is a review of our consideration of topics essential to an understanding of lightning upset. Section II is a first cut at integrating the knowledge gained in Sec. I into an approach to comparing one design against another. Section III addresses an approach to lightning-induced transient fault effects in digital systems from the fault-tolerant aspects of the problem. Section IV is a list of tasks which require further work.

I. Aspects of Assessing Lightning Upset

There are many aspects to assessing the lightning upset potential of an airborne digital system. The overall items which will be addressed are:

A. Characterizing the electromagnetic environment of the digital system, and

B. Consideration of system properties and system effects.

A. Electromagnetic Environment

Any upset-potential assessment techniques must take into account the electromagnetic environment in which the system is to be operated. The upset potential of any particular system against transients generated by lightning will depend on whetner single or multiple faults are generated within one or several units. While it is relatively simple to design a fault-tolerant system which will deal with single faults in a single unit, it becomes much more difficult to design a system which deals with multiple faults in several units simultaneously. The upset potential of a more complex system will also be mcre difficult to assess. It is therefore beneficial to examine the electromagnetic environment and to make it as benign as is practical.

Research efforts to describe the electrical transients created by lightning in an aircraft environment have been carried out for several years (Reference A). The lightning stroke, either nearby or direct, interacts with the airframe and produces the stress of the computer system inside. Analysis and measurement

can determine the threshold of the system for upset (or damage, or any other criterion). In order for the system to function according to the criterion selected we must insure that the stress at the system level is below the threshold of the system. However, it may not be practical to fulfill this condition for the entire system, in particular if the system is a distributed one; or it may not be practical to design the system for a worst-case lightning strike.

1. Lightning Stress

When lightning strikes or occurs near an aircraft, intense electromagnetic (EM) fields exist outside the aircraft which create large currents and voltages on the aircraft exterior. Since the aircraft exterior is not a perfect shield, these large currents and voltages on the aircraft exterior create large EM fields inside the aircraft. The internal EM fields induce current and voltage transients on the internal wires. The lightning induced current and voltage transients are conducted by the wires to the inputs of integrated circuits (ICs). The ICs connected to wires on which lightning-induced transients exist can be upset. An important aspect of the problem is that the lightning-induced transients exist simultaneously on many wires connected to many ICs. The effects caused upon ICs by the simultaneous occurrence of a large number of lightning-induced transients needs to be investigated. The investigation should consider both individual ICs of various complexity and small subsystems of ICs such as a single printed circuit board. The investigation should include a

study phase in which related EMP and other EMC investigations are reviewed carefully. Then appropriate analytical methods, experiments, and computer simulations should be proposed to determine the new lightning-induced upset information believed needed. How the new lightning-induced upset information for individual ICs or small collection of ICs is to be used to assess overall system upset effects should be considered throughout the effort.

2. Shielding

In order to obtain a useable electromagnetic environment for each unit of a system, i.e., to avoid unmanageable stress, the design of the system should incorporate the concepts of system topology. These concepts were originally developed in connection with hardening of ground-based facilities against nuclear EMP, and also in connection with electromagnetic interference control (Reference B,C). The topological model leads to the design of a cost-effective system whose reliability can be predicted with confidence. The basic ideas of the topological zoning concept will be summarized in the next paragraph, details can be found in the two references cited.

In the topological view, a system is decomposed into different "zone" or "volumes". Each zone is separated from the neighboring zone by a barrier which is substantially impervious to electromagnetic waves or conducted currents. Although the ideal barrier is a perfectly conducting closed shield, this ideal need not be achieved in practice, because any metal shield of structural thickness offers sufficient attenuation against diffusion of

electromagnetic fields. The most serious violation of the requirement that the shield be closed are penetrations by insulated conductors. These penetrations are of course necessary but they must be treated at the barrier with devices like filters, surge arrestors, limiters, etc. These devices attempt to close the shield, and they can effectively achieve this outside the frequency band used for data transmission or above a maximum voltage level. The system design should begin with an identification of the boundaries of the various zones. If the boundaries coincide with metal shields, then the configuration control consists simply of identifying the penetrations and treating them to the extent possible. (Apertures are usually much less important, but they should also be examined, see Reference B and C.) If the system and the barriers are designed such that each barrier attenuates external noise to a level below internal noise* generated by the subsystem itself, then an additional advantage is gained: the system continually tests itself against failure because any outside transient will be reduced to a transient no larger than the ones generated by the system itself. It may not be practical to require that much attenuation of a particular barrier; a tradeoff between isolation requirements and upset potential must be performed. Note that a ground conductor should never penetrate a shield because it compromises the integrity of that boundary.

How are the zones identified in an aircraft? To be cost effective it is essential to take advantage of existing

---

*Internal noise here does not mean the steady-state background but rather the maximum transient generated by the system.

boundaries. For example, in an aircraft with a metal fuselage the first natural boundary is the fuselage itself. The skin of the aircraft will dramatically reduce the amplitude of the external lightning transient. The next boundary may be an equipment rack, or an equipment box, and so on. Thus each successive zone becomes electromagnetically quieter than the one outside of it. In an aircraft with a fuselage made of advanced composite material the first barrier may be the equipment rack. Cable trays or conduit connecting different racks could ensure the integrity of this first level of shielding. In some cases, especially when two units are separated by a large distance, fiber optics could be used to save weight.

It may turn out that a system cannot be designed economically such that every procesor is located in a quiet zone, especially in an advanced composite aircraft. In such a case, high quality shielding could be applied around a compact "inner sanctum". Except for power lines (properly treated with filters and surge arrestors) there would be no electrical penetration of the innermost barrier. All data lines would consist of fiber optics. Such a design would have the advantage that the proper functioning of a processor in the inner sanctum could be relied on with very high confidence. This unit would be unaffected by lightning or other transients and could therefore be used to do the necessary checking of peripheral units to determine their proper functioning.

B. System Effects of Upset

As discussed above, computer systems, and more specifically

fault-tolerant computer systems in aircraft, can be subject to harsh electromagnetic environments. There is a need to assess system designs, both hardware and software, for susceptibility to upset.

1. Stress Characterization

At a system level the currently foreseen stress of upset is the disruption of vulnerable registers and memory, and the disruption of clocks. A well designed system will eventually restore the disrupted registers and memories, and will restore the proper phase relations among its clocks. Obviously the system cannot begin to recover until the stress ends. From knowledge of the electromagnetic environment as attenuated through shielding (which is a part of the system design) one must characterize the maximum disruption that a design is intended to withstand. As a rough cut, this can be characterized by two parameters TV and TR defined as follows:

1. TV is maximum duration of electromagnetic burst above the the noise immunity threshold of the system.
2. TR is the minimum time between electromagnetic bursts.

Because lightning often comes not in single bursts, but in multiple bursts, a more complete characterization is needed. This is a subject of further research.

By fault-tolerant designs we mean designs that involve spatial redundancy--usually with cross-checking and provision for reconfiguration in order to avoid system malfunction in the face of component failures. In effect, "fault-tolerant" means "intel-

ligently redundant". Redundancy is a powerful approach to component failures that are sparse (as opposed to numerous over time), hard (as opposed to marginal), and enduring. Upset results in error conditions that are not sparse, but rather can be dense, many or even all copies of a component can be effected at once.

The hope for overcoming upset is that the environmental effect which causes it is short-lived.

2. System Characterization

Two aspects of system response to the effects of upset need to be considered. One is the response due to design features not including features present on behalf of fault-tolerance. The other aspect is system response.

It is convenient to think of a redundant design as acting like a non-redundant equivalent machine with extraordinarily reliable components. In this way one partitions the analysis of the response of a fault-tolerant (i.e., redundant) system into two parts:

1) Determine the response of the "equivalent" non-redundant "system" to lightning upset.
2) Set aside the equivalent system and consider the actual fault-tolerant design, and determine the coupling between design fractures present on behalf of fault-tolerance and the other features in this joint response to the stress of upset.

Part 1) is much easier and system designs which are not satisfactory with respect to that can be ruled out without the

necessity of undertaking part 2).

Concept of Flush-time

Some process control systems or subsystems can be designed to have what is called a flush time. The basic model is that of the controller for a washing machine. For a flush time to exist the subsystem must refresh all writable memory from input data, or calculations based on input data within a fixed time - the flush time. A subsystem with a flush time will recover its proper operation within the flush time after upset stress has terminated.

Correlated Failures in Redundant Systems

By fault-tolerant designs we mean designs that involve spatial redundancy -- usually with cross-checking and provision for reconfiguration in order to avoid system malfunction in the face of component failures. In effect, "fault-tolerant" means "intelligently redundant". Redundancy is a powerful approach to component failures that are sparse (as opposed to numerous over time), hard (as opposed to marginal), and enduring. Upset results in error conditions that are not sparse, but rather can be dense, many or even all copies of a component can be effected at once. The hope for overcoming upset is that the environmental effect which causes it is short-lived.

Fault tolerant computer systems have for the most part been designed under a basic assumption that their environment is random and that failures occur in a random uncorrelated fashion. This assumption may not be valid in the case of lightning induced upset

of an avionics system. First, it is not clear how the electromagnetic field will develop internal to the aircraft and it is not clear what effect the physical arrangement of a distributed avionics system will have. So until real time experience and data is obtained it is best to assume that there may be positive correlations between failures. There seems to be a need for basic theoretical and experimental research in this area. At the most elementary level the problem is simply how to construct an adequate voting system among some small number computers whose correlation matrix is known. This problem is itself not totally trivial even when all of the correlations are zero.

In a distributed system failures may occur simultaneously in physically separated components due to, for example, leads that are in close physical proximity to paths of lightning induced currents. It is simply not clear at this time what form the appropriate counter measures should take. Basic experimental data for metal and composite bodied aircraft in lightning strike is needed.

It will not be possible to completely decouple the avionics system. However, part of the design strategy should be to make various subsystems as independent as possible. On the other hand, effective monitoring should require that the systems be linked as closely as possible. Thus basic tradeoffs must be made between minimizing correlations and effective redundancy.

Basic research is needed for the design of fault tolerant computers in the presence of correlated failures. A promising approach would be to adaptivly identify the correlation matrix and to adjust the probability that an individual machine is faulty as

failures occur. Other approaches should be given careful attention.

3. Upset Detection

Given (i) generic models of lightning-caused pulses which must be dealt with in the sense that these pulses can potentially penetrate the established barriers and (ii) an understanding of the degree to which isolated/well-protected barriers can be established (i.e., quiet zones) together with their associated costs, research efforts should be be directed at the following two issues. First, efforts should be aimed at developing new types of (what we will refer to in the following as) upset detection mechanisms/schemes that are fundamentally different from those now used for single, permanent failure-mode detection. The term "upset detection mechanism/scheme" is used here to broadly describe any software or hardware or firmware approach (or combination of these) to the testing/monitoring of operational aspects of a digital system. The specific details or properties of an upset detection mechanism/scheme are, of course, that upon which research will focus.

Regardless of the particular realization of an upset detection mechanism/scheme, they will, however, all be highly specialized or sensitized to the types of upsets that can result in a digital system as determined by (i) and (ii) above. Of fundamental significance will be the requirement that any useful, realistic upset detection mechanism/scheme must itself be highly protected from upsets. This means that such upset detection

mechanism/schemes must be located in highly shielded zones. Because of the expense of such zones, these upset detection mechanisms/schemes should then also be "small" relative to the overall system. (Otherwise, why not protect the whole system in such a quiet zone?) It is reasonable to suggest that the "smallness" of a detection mechanism/scheme is closely related to the number of failures and subtleness of the failures for which it is responsible. This then leads to the second issue at which research should be directed. Again given (i) and (ii) above, the requirements on the initial system design (hardware and software) that would "force" or "restrict" the upsets to take on more controlled and observable forms should be investigated. In other words, can the system be designed such that by fully exploiting (i) and (ii) above, the class of upsets that must be handled is limited to a finite overt set? This is another way of saying that if you cannot avoid the upsets, then learn to live with them by designing the system such that upsets take the form of deviant, but nevertheless clearly observable systems operation. By exploiting this forced observability, the upsets can potentially be studied experimentally. These studies will lead to models of upsets based upon the more salient features of the experimentally observed phenomena. Such experimentation will require the development of innovative fault injection strategies, and this in turn will require the availability of state-of-the-art laboratory instrumentation. The expense of establishing a number of such laboratory facilities would be well worth the possible outcomes. In particular, the possibility of emulating upsets to evaluate

detection mechanisms/schemes can only be addressed after injection experimentation on a variety of prototype systems has taken place.

It should be pointed out that the consideration of upsets being suggested here is in contrast to reduncancy voting designs wherein the details of a single upset do not matter, as the effects of an upset on data output are simply voted away. In the approach being proposed here, the details of the upsets are of the essence, for it is these details/features that allow us to consider highly specialized, but small and well-isolated upset detection mechanisms and schemes.

One final research area related to the above two issues is the possibility of developing "warning sensors" that trigger the upset detection mechanisms/schemes relative to the current state of operation of the system. A "warning sensor" would be a circuit that was simply meant to flag a "high probability of upset condition". The isolated/well-protected upset detection mechanism/scheme could then take specialized actions dependent upon the current system state. The upset detection mechanisms/schemes would then be conditioned by system state in the sense that in a certain state of operation, an upset detection mechanism/scheme when alerted of a highly probable upset condition by a warning sensor could turn its attention to, say, feature A, of that which it is monitoring whereas if the system were in a different state when the warning was received, the upset detection mechanism/scheme might turn its attention to, say, feature B. This suggets that the upset detection mechanism/scheme is somewhat intelligent background monitor.

Differentiation of External Environment from Computer Failure

The phenomena of lightning-induced upset presents special problems in the detection of system failure and determining if a real failure has occurred or is only perceived. Lightning induced upset occurs most probably in areas of high atmospheric turbulence which greatly stresses the airframe components. There may be large excursions of dynamic pressures and significant altitude and velocity changes may be expected. If the detection of a computer malfunction is left to the pilot two things can occur. First, he may not notice because of an already heavy work load. But worse he may misinterpret information he is receiving. He might interpret the results of an actual computer failure as being due to the external environment and not take appropriate corrective measures. Or he might interpret the influence of the external invironment as being caused by a computer failure. It is precisely at this time that the computer controlled avionics system is most important. If he attempts to reconfigure or to take manual control, catastrophic failure could result (especially in a control configured aircraft such as those built of advanced composite materials). The pilot is used here only for analogy, for any master detection scheme will have precisely these same problems.

The problem is basically the classical problem in communications of a signal in a noisy environment. In this case the signal we want to detect contains the information that a failure has occured. If the noise to signal ratio is high the detection process may be very difficult if not impossible.

A successful design then should incorporate some of the following features. One, the detection of a computer failure should be made as obvious as possible. This clearly involves a tradeoff in design between making a failure have no effect and making it readily detectable. Two, there should be dynamic cross checking between different sections of the airframe to determine the true dynamic parameters of aircraft. Unfortunately this creates a need for significant amounts of additional computational capability and leads to an increased probability of computer failure. Design tradeoffs will be necessary.

The other possibility mentioned was that there could be actual mechanical failures induced by the turbulance. Any state of the art avionics system should contain a failure detection and diagnostics subroutine. Great care should be taken in the design to allow for the differentiation of frame failures and computer failures.

4. Lists of positive and negative features

Several aspects of a system design can be examined by methods that include checking for the existence of known essential features, and using various modeling techniques to indicate something about expected consequences of any troublesome features that are found. The only methods applicable to designs are methods of analyzing as opposed to test. Analytic methods are based on a list of necessary or at least positive features and a list of negative features. The most positive statement about a design that can result from any method of design analysis is that the design exhi-

bits all the listed necessary features and none of the listed troublesome features. This does not mean that a system built to the design will behave well when upset. In other words, methods can be developed that can exclude known dumb featues and can demand known necessary features. Potentially, design analysis can act to screen out obviously weak systems, and to suggest gross areas for their correction. The objectives of the work reported and proposed have to develop this potential.

Analysis is based on lists of features that are known or imagined as possibilities - features of the design and features of the environment. Experience produces surprises which modify these features. Thus analysis can help to avoid the avoidable but cannot substitute for experience.

Somewhat Obvious Design Practices Which Improve the Lightning Upset-Tolerance of a Computing System

The design of a completely upset-tolerant computing system is an insurmountable problem. However, it is possible to design a system in such a way as to include a certain degree of upset-tolerance. By examining the various problems caused by upset and considering some of the most obvious design remedies, it is possible to improve the upset-tolerance of certain computer system designs by a large degree.

The first and probably most obvious means of providing a degree of upset-tolerance is to store all programs in read-only memory (ROM). The information stored therein is non-volatile and thus not susceptible to change when under the influence of an

upset phenomenon.

Avoiding the use of multiple byte or multiple word instructions in the design of a computer system is another means of improving upset-tolerance. Realizing that an upset phenomenon can cause a random change in program counter contents, it can be seen that the use of multiple-segment instructions allows for the possibility of reading an incorrect or invalid instruction from memory. This occurs because an upset may not leave the program counter pointing to a location which is not an intended instruction boundary. After this type of occurrence it is very difficult to re-synchronize the program execution.

Definition and/or implementation of all possible instruction codes is another means of improving upset-tolerance. This point is particularly important to instruction code designs which are necessarily multiple-byte or multiple-word. By defining all possible codes an appropriate action can be taken when non-useful (otherwise invalid) instruction codes are fetched. (Fetching of invalid instructions is highly likely to occur when the program counter is upset and left pointing to an instruction boundary.

In the design of the computer systems control unit sequencer, it is important that all possible state transitions be defined. In many cases control sequencer designs are minimized by taking advantage of the fact that certain inputs are never expected or that certain states are never expected to be reached. This allows the designer to use the don't care concept in filling out a next state table. It becomes fairly obvious when considering lightning upset that the control sequencer inputs and states may become

unpredictable in the presence of noise. Considering this it is necessary to be specific about all next state transitions in the control sequencer in order that the system not get "hung up" in an undefined control loop.

These are but a few of the obvious "good" design principles which should be included in an upset tolerant system. These principles are similar to those which might be included in the design of a system which is to be fault-tolerant in the face of component failure. However, there is considerable difference between upset-tolerance and fault-tolerance and this difference must be thoroughly considered in order that an appropriate set of "obvious" design principles for upset-tolerant digital systems may be specified.

II. First Cut of a Theory of Comparing One Design with Another for Susceptibility to Lightning Upset

A method is needed to compare the relative strengths and weaknesses of two candidate designs for fault-tolerant avionic computers, with respect to their susceptibility to lightning upset. This involves comparative analysis of the two designs. This means that one analyzes how systems built according to different designs would behave when subjected to some test or tests. As a preliminary to discussing methods of analysis and simulation, we address the issue of comparative tests. In Sec. A we pretend that the designs have beem implemented, and that computer hardware built according to those designs is available for testing. The issue is the design of tests to compare a computer built to one design against a computer built to another design, both operating in the presence of lightning. In Sec. B we turn to the issue of determining, in so far as one can, what the results of such a test would be, without actually running the test or having the hardware.

A. Tests of Comparative Susceptibility

It is instructive to think about the requirements for such a comparator. The first is that the comparator connectors and the computer connectors have to be compatible, and that the format of the signal transmissions also be compatible. A reasonable approach is to specify a mock aircraft in which the computer is to operate. In particular, one specifies an I/O scheme which both designs must adhere to. Thus a computer of any acceptible design

will be supposed to write into standard output registers.

A second requirement is that the lightning stress imposed on the computers has to be fair, in the sense that at least on the average the computers have to be equally stressed. Further research is required to clarify what 'equally stressed' means when applied to computers of different design.

Thirdly, there must be a means of comparing upset. This is different from comparing two computers; to compare upset, one needs a two-level scheme of comparison. The first level is to compare a computer of one design. which is subject to upset, against a computer of the same design, which is not upset. This comparison is necessary to measure upset at all. Then the upset for one design must be compared with the upset for the other design. Hence the comparator scheme must be elaborated: two computers of design 1, two computers of design 2, and three comparators must be used, as shown in Fig. 1. Two computers of design 1, one subject to upset and the other protected from upset, are connected to a first comparator. The third (i.e., output) port of this comparator reports on the upset experienced by the computer of design 1. A similar set up is used to produce a report on the third port of the second comparator; this report is of the upset experienced by the computer of design 2. The outputs of both comparators are connected to the inputs of a third comparator, and the third port of the third comparator reports on the relative upset of one design vs. the other design.

A fourth requirment is for measures of upset, to be used by a comparator that compares an upset computer with a computer of the

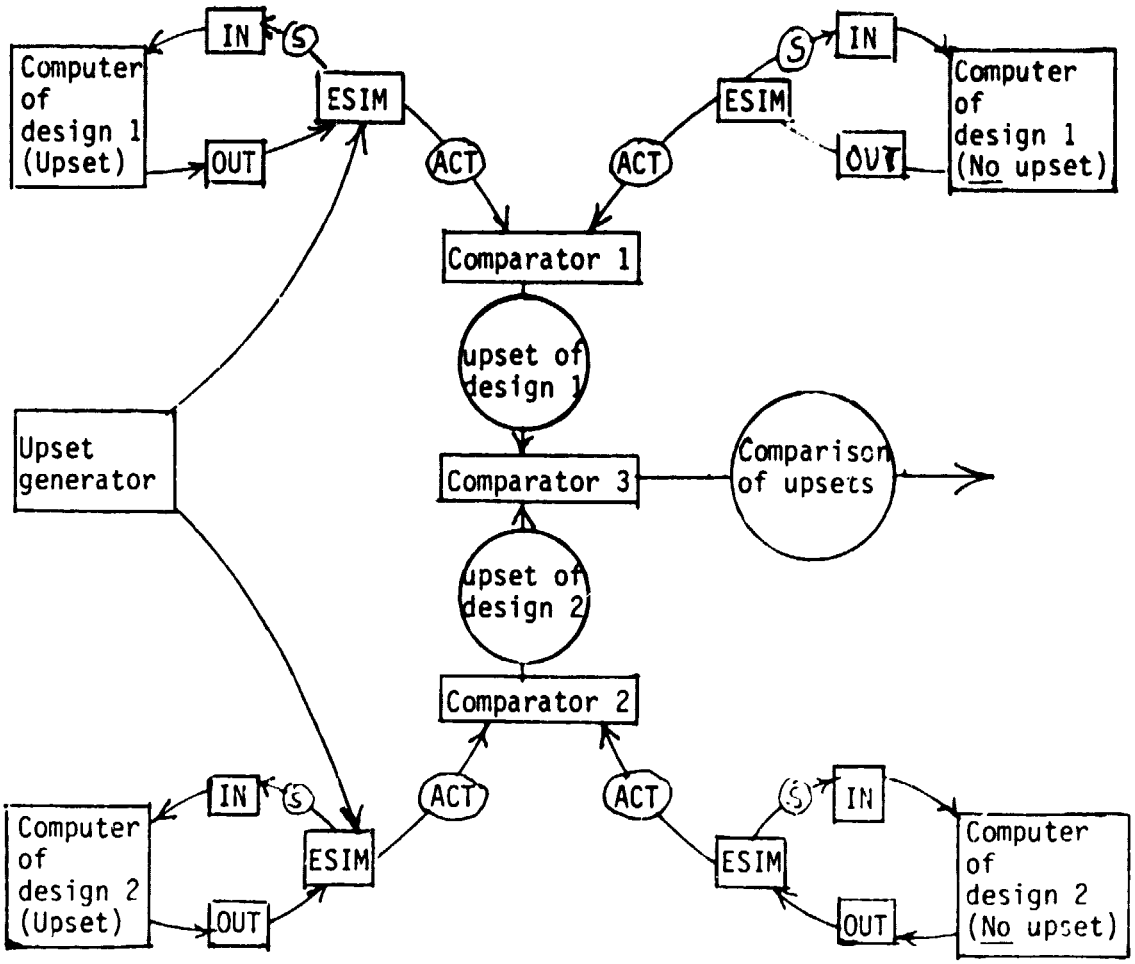

Legend: ACT = simulated control voltages to actuators;
S = simulated sensor reports;
IN = input registers of computer;
OUT = output registers of computer;
ESIM = simulation of mock aircraft, including behavior of communications network and flight.

Figure 1: Test set-up for comparing computers for susceptibility to upset.

same design which is not upset. As a starting point, one can think of the output registers as acting on and through a communications network to direct actuators. We assume that both designs function through identically designed networks; these networks are part of the mock aircraft. It is then possible to analyze the effect of irregularities in output on the performance of the aircraft. In other words, the measure of upset will be some function of the computer output. More precisely, it will be a function of the outputs of the computer which is upset and the computer of the same design which is not upset. A simple comparison, integrated over time, would give some information. A better measure would allow for the conversion of digital signals to analog signals, would consider the response time of the actuators, would account for possible influence of some registers on the routing of signals from other registers, and would weigh the comparision for the critically of the control function affected.

There are two modes of comparison, an open-loop mode and a closed-loop mode. In the open-loop mode the input registers (which convey data to the computer) convey the same data to each computer, independent of what the computer puts in the output registers. The closed loop mode of comparison includes a way (e.g. simulation) to allow for the feedback through the environment i.e. if the computer generates a dive of the aircraft, input registers will bring back sensor responses to that dive. In this mode two computers which operate differently in response to upset will receive different input signals. Both modes of comparison are of interest.

The comparison of the outputs of two computers can be wildly misleading if the two are not exactly synchronized. Because fault-tolerant computers involve independent clocks on each redundant module, and are required to be diverted for fault handling, and because some random uptset phenomena will trigger fault-handling action, the upset computer will not be in synchrony with the computer which is not upset. For this reason the comparator should not be attached to the computer outputs, but should be attached to the control voltages of the actuators. It is unlikely that in a test one would have actually implemented the mock aircraft, but a simulation is required, so that the control voltages on actuators that respond to the computers under test are available. The closed-loop test scheme is diagrammed in Fig. 1; for open-loop testing the input registers of both computers of a given design are driven by the non-upset environmental simulation.

B. Analysis and Simulation of Comparative Susceptibility

We now consider what can be learned from a detailed description of the computer designs, without having the computers on hand. In a nut shell, one needs to be able to analyze the affect of upsetting the input registers, and possibly the clocks, on the output registers. If normal practice is followed in the design of the hardware and the operating system, such analysis will not be possible. On the other hand, it will not be needed, because normal practice results in such bad performance under upset that the design can be ruled out merely for being "normal". Thus a primary design requirement whould be that the design can be analyzed for

its response to any sequence of inputs and any noise condition on its clock lines. Section I.B.4 addresses a few of the features that facilitate such analysis; further research is needed to examine the issue of analyzeability in more detail.

III. An Approach to Lightning-Induced Transient Fault Effects in Digital Systems

A digital system accepts inputs, and acting on these inputs, in conjunction with previous inputs, produces an output. Figure 2 shows such a system considered separate from its environment. However, to consider lightning-induced transient faults, a digital system is better envisioned as one element of the overall operating environment. In normal operation, the operating environment generates a class of input signals to the digital sytem. The outputs of the system, through actuators, alter the environment, and can thereby affect the inputs to the control system. This feedback is shown in Figure 3. This conventional model is based on the fault-free assumption.

The presence of lightning-induced transient faults can alter the functional effect of the digital system, and perturb the environment so that inputs are no longer the same as in the fault-free case. Shown in Figure 4, the effect of faults on the system is as if an additonal set of inputs to the system were present, defining the transfer function between the conventional system inputs and outputs. The fault-free condition would be but one case of this abstract, fault transformable control system, corresponding to one set of inputs present at the function transform input. This enlarged system (the expanded equivalent digital system), which models all possible functions of the original system presented with faults, need only be considered in the fault-free case; by application of function transform inputs, it operates fault-free in the same way as the original system subjected to any specified

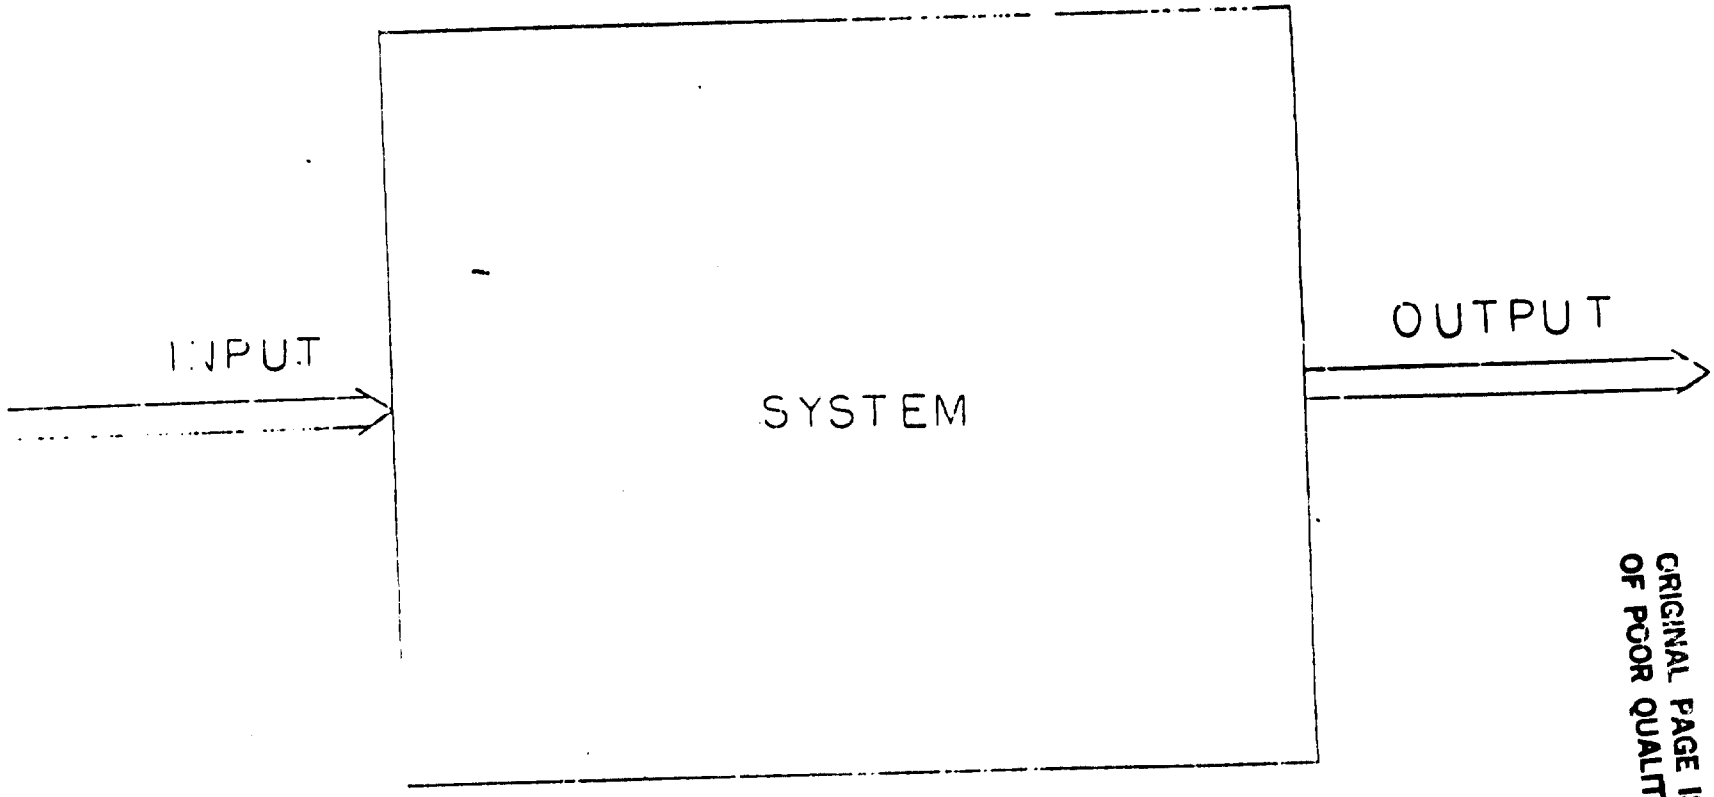

FIGURE 2

SYSTEM CONSIDERED ALONE

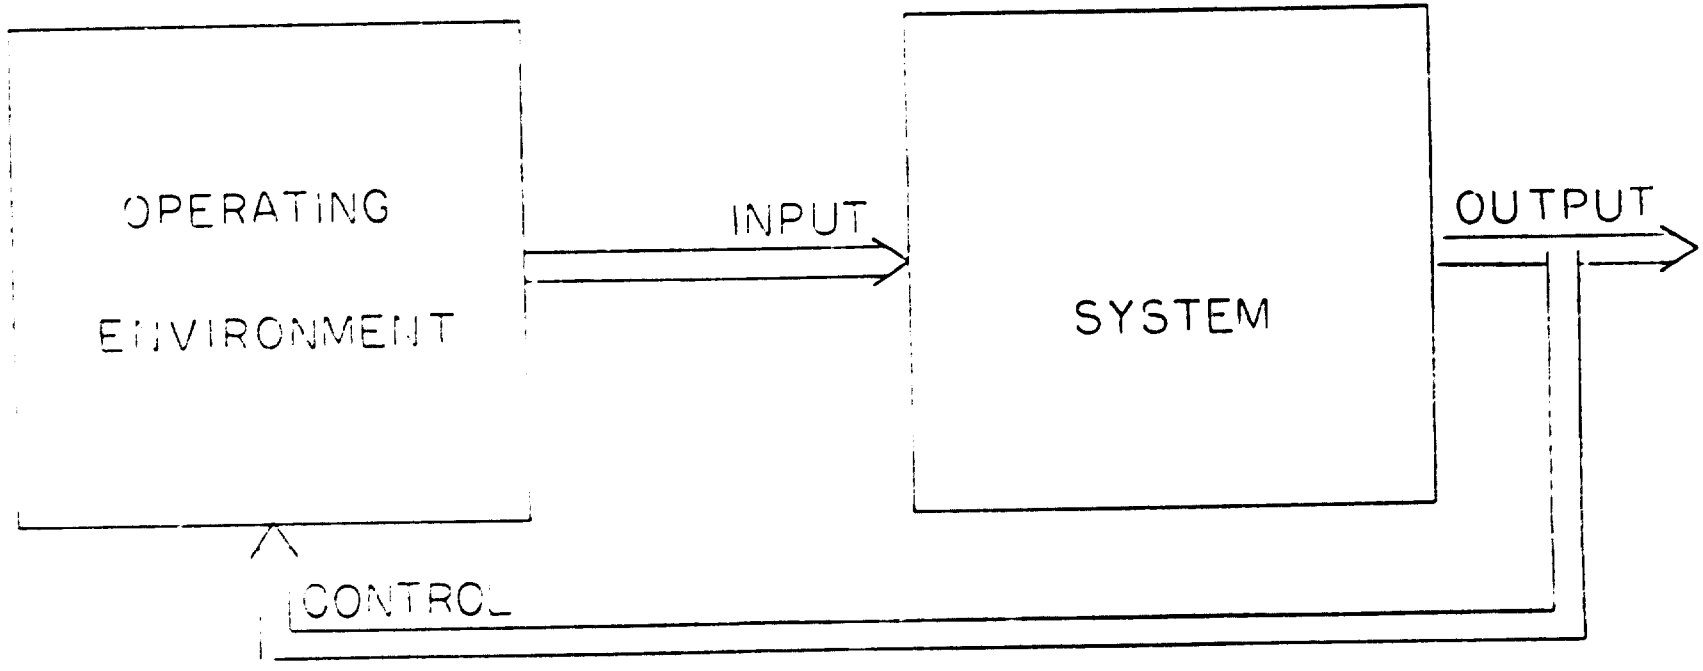

FIGURE 3.

SYSTEM SITUATED IN ITS ENVIRONMENT

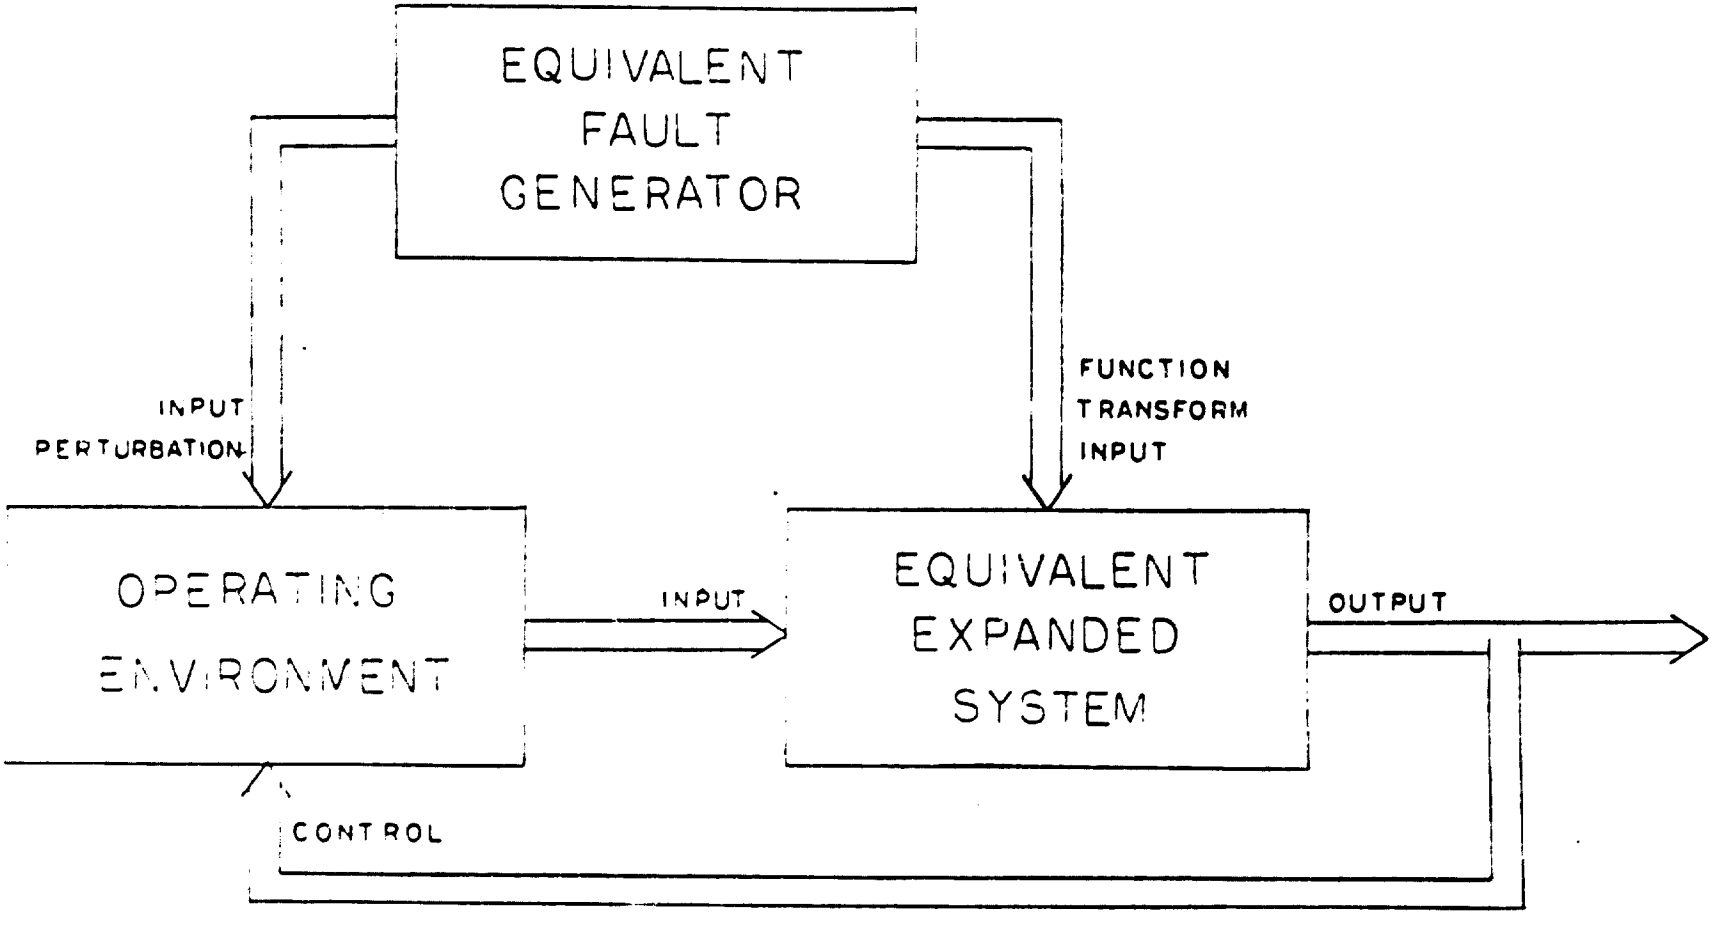

FIGURE 4

SYSTEM SITUATED IN EQUIVALENT HOSTILE ENVIRONMENT

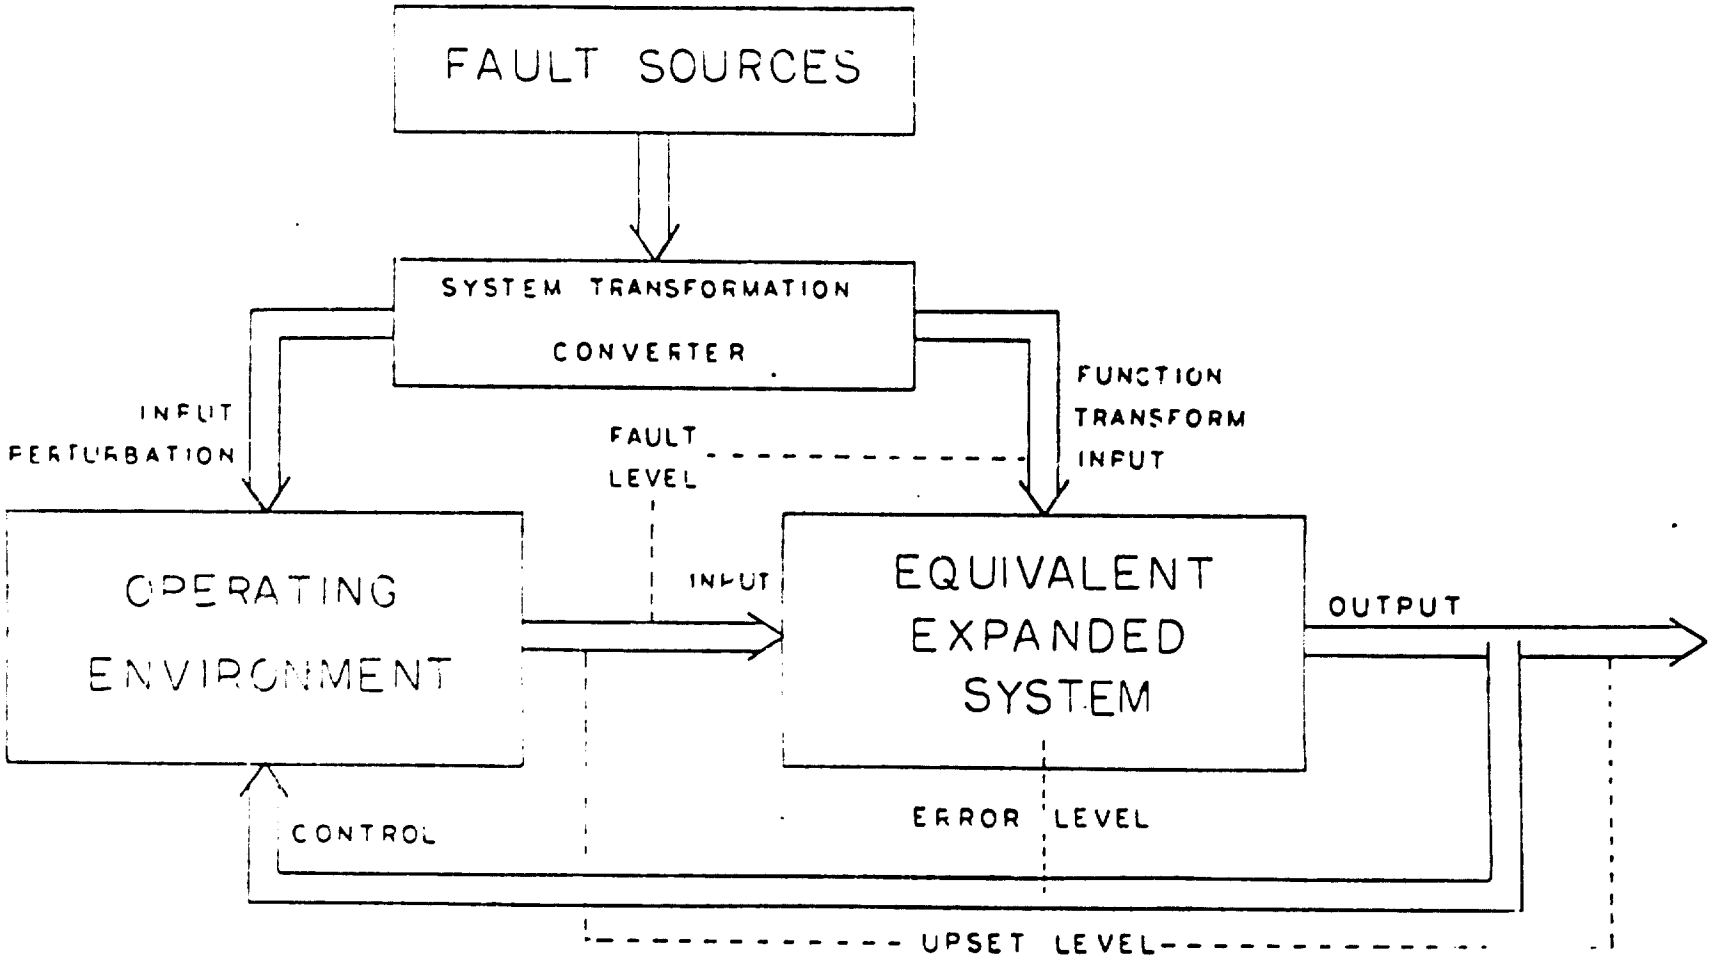

FIGURE 5

SYSTEM SITUATED IN HOSTILE ENVIRONMENT

faults. An equivalent fault generator produces digital outputs for the expanded digital system which correspond with the actual faults occuring in the hostile environment. Figure 5 shows this equivalent fault generator broken into two parts: the fault sources, and the system transformation converter. The converter accepts lightning generated analog fault inputs at various points in the circuitry, and produces the digital function transformaton outputs needed for the expanded digital system to behave as the original did in the presence of faults.

This abstract model contains portions which must be determined by an intensive research program. Given complete implementation information for the original system, the expanded digital system can be determined through a thorough fault analysis of the original digital system. For complex digital equipment, this will be a very difficult task; merely functional, or even logical, specificatons are not sufficient for fault analysis, and chip level information can seldom be obtained for modern integrated circuits. Nevertheless, there is a reasonable possibility that this can be practically done. However, the real issue here is the system transformation converter. It is the junction between a vast body of knowledge concerning detailed fault sources, and a large, but theoretically well understood, fault-free digital system. Very little is known about this interface -- therefore, a research program should be supported which will take this model from only a basis for abstract understanding to a directly implementable approach. Issues regarding faults, errors, and upsets that are pertinent to such work will be disscussed in the following.

A fault is a logical difference at the site of a circuit failure between faulty and fault-free devices physical failure, and is expected to make transitions between active and inactive states during the lifetime of the equipment. A lightning-induced transient fault is not due to a circuit failure, but an environmental condition which the equipment was not designed to tolerate. Information at this fault source level is not digital, but analog in nature. For example, the actual fault source can be an internal logical signal modification such that it satisfies neither the logic high or low digital requirements. In such a case, it is not possible, in general, to predict how a digital circuit will react. Moreover, while the digital circuit is a clocked, synchronous digital system, the actual lightning-induced faults caused are seldom "well-behaved" in the sense that the fault is only synchronously active or inactive with system clock. The well-behaved assumption greatly simplifies predictions of the fault's operational implications, but clearly does not reflect real situations.

In Figure 5, the fault sources provide analog inputs to the system transformation converter. The output of the converer is the first point where the fault sources produce a digital output which is clearly different from the fault-free condition, and threfore this is the fault level. Theoretically, knowledge of the signals at the fault level is sufficient to predict the behavior of the system in the presence of faults; determining the faults then amounts to identifying the system transformation converter. Physical characteristics of the digital system contribute to the makeup of the converter, so "fault models" which utilize fault

source information exclusively, ignoring the system being subjected to those fault sources, are not sufficient to identify the converter, and therefore not sufficient to model the overall fault/system interaction.

An incorrect logic value at a fault site propagates to other parts of the circuit. These logic differences between faulty and fault-free systems are called errors, with the implication that an error at a failure site is called a fault. Errors can be considered on all lines of a circuit; in Figure 5, they appear at the digital system outputs, and on all differing lines internal to the digital system itself.

It has been seen that the fault level is not useful as a description of fault/system interaction in complex digital circuits, because observation is not possible at this level. The error level is one step removed from the fault level. A lightning-induced transient fault can cause a state change in the digital system. A continuous string of errors can result if the faulty circuit is not forced into the correct state, that of the fault-free circuit. Simple loss of synchronization will yield endless errors, after the I/T fault has disappeared.

In modern digital systems, consisting of LSI integrated circuits, all circuit lines are not accessible; only pins of the packages are obseervable. In the attempt to find the lowest observable level from which to view faults, a complete description of errors is not usable, since most of these observation points are not available. A subset of circuit lines are available, but, even accepting error propagation latency, the conventional defini-

tion of error suffers from the loss of synchronization problem. To make errors a useful point of observation, a new definition is made such that logic signal differences at the observation points between faulty and fault-free systems will only be interpreted as errors when the digital system is actually being driven by a fault. If a set of logic signal differences are observed over one clock cycle, then an error will be noted over the next clock cycle only if there are logic signal differences at the observation points between the faulty circuit and a fault-free circuit which has been forced into the exact same state as that of the faulty circuit at the end of the previous clock cycle. With this interpretation of an error, an error of burst length b occurs if there are logic signal differences at the observation points over b clock cycles. Hence, when an error of burst length b is observed, the circuit has been driven by the actual fault source for b consecutive clock cycles, less the latency time. It can be argued that this definition of burst length errors is the lowest observation level which is of any value in modern digital systems. However, as will now be discussed, the upset level is a more useful perspective for the goals of the research program being proposed.

A digital system is designed to perform some function. The primary concern is how faults affect the performance of the design function. The presence of faults perturbs the system, or "upsets" it. The viewpoint; which observes fault effect at this higher, functional level, will be referred to as the upset level. The transfer function described by the system outputs, shown in Figure

5, in relation to the signal inputs, is the observation point for the upset level.

At the upset level, a system is viewed as responding to the arrival and departure of a transient fault in two states -- the same as standard system theory separates a general system's response to any input. The transient fault is an input to the expanded equivalent system of Figure 4 and 5. This system's response to the transient fault input will be the same as any other system's response to an input: there will be components both of the transient response and the steady state response. The difference here is that with standard system theory, the steady response is due to a driving input, which remains, and the transient response is due to the system's response to the input change. In the case of the expanded equivalent system with regard to its response to transient fault inputs, the transient fault arrives and then disappears. In the purest sense, then, the system' response would be composed entirely of the transient response, with no steady state response. However, for this digital system driven with transient fault inputs, state changes can cause lasting effects on the system after the departure of the transient fault. These effects must be classified as steady state effects at the upset level. Fault effects present during, and shortly after, the time when the system is being driven by the transient fault will appropriately be called transient effects. (The word "transient" here is used differently from that in the term "transient fault." There can be both "transient" and "steady state" responses to a single "transient fault." The single word "transient" is retained

in this discussion for each of these separate meanings because both are in agreement with standard usages.)

It is possible to observe transient and steady state effects of transient faults because of the functional interpretation of the upset level. Typically, transient responses are incorrect data values, loss of time or synchronization, or skipping a computation step. For example, consider a digital system behaving partially as a moving average filter; the outputs depend primarily on current and recently applied inputs. Effects of an input disappear with time. For this type of digital system, transient output perturbations due to transient fault will also vanish over a period of time. This will be true only if the system function remains unchanged by the fault.

Steady state responses are functional transformations. After the departure of the transient fault, the system is no longer performing the same transfer function between its signal inputs and outputs as prior to the transient fault arrival. For the moving average filter example, a steady state response to a transient fault modifies the filter algorithm. Since the filter may no longer be a moving average type at all, there can be no expectation that output perturbations will disappear with time; in fact, with a function change, the fault effects will not disappear.

A more sophisticated digital system could monitor internal states and external events to determine if transient faults have caused a transient effect on data processing, and initiate recovery procedures when detected. Transient system responses can be tolerated in this fashion. Steady state system responses to a

transient fault would transfer execution from the algorithm being performed, after which there would be no resonable hope for system recovery. Because of the relatively drastic consequences of a single transient fault that the steady state fault response can show at the upsest level, as compared with the transient response, probably the steady state effects of transient faults should first be considered in this research program.

To concentrate on the steady state fault response of a digital system, the system's functional I/O relationships must be characterized. The operation of the system can be completely described in terms of a finite set of mutually exclusive functional states, covering all possible transfer functions of the system in Figure 2. This set of functional states will be referred to as the containment set. All possible system states must cause functional operation of one of the elements of the containment set. This set must include all possible valid functional states of the fault-free system (that of Figure 1), but, this set must also include invalid functional states, not explicitly designed into the system but into which the system, can nevertheless be driven by a transient fault. Arrival at these states, of course, can only be through those additonal functional states created by nonzero fault inputs to the expanded equivalent system of Figure 5. (An imporperly designed system may also arrive at these states due to unexpected inputs.)

As was previously discussed, an understanding of system responses to particular fault sourcs requires a detailed knowledge of the system transformation converter of Figure 5, and this appears

to be a highly worthy research goal that should be supported. However, when dealing exclusively with transient faults, and when ignoring transient effects of these faults on the system, the steady state effects can be described in terms of the containment set; and this set is completely defined by the structure of the expanded equivalent system. Moreover, since the steady state fault response is of primary concern (after the departure of the transient fault) the fault outputs of the system transformation converter will all be zero; and the expanded equivalent system is reduced to the original system. This means that the containment set is completely described by the original digital system itself; no fault information or system expansion is needed to obtain the containment set. This advantageous condition is a result of not considering solid (permanent) faults, and then ignoring the transient effects of the remaining faults. Nevertheless, this containment set is the basis for the analysis of lasting effects of transient faults, and the likelihood of practically determining containment sets can be seen to be promising. It is not at all clear that a finite containment set exists for all digital systems; and this is another important research area. It must be determined what type of system modifications are necessary to produce a useful, finite containment set, and whether or not these modifications are overly restrictive.

Hence, we have proposed a new approach to lightning-induced transient faults in digital systems. This approach has many major questions connected with it. It could be the basis of a number of important research investigations.

IV. Tasks for Future Work

A. Establishment of electromagnetic zone definitions for airborne computer systems. Initial determination of how these zones can be implemented in airborne systems.

B. Formulate methods which can be used to determine the EM thresholds which when exceeded will cause detectable changes in digital electronic components.

C. Determine the amount of multiplicity present in lightning upset by considering the electronic effects on all circuits within a certain electromagnetic zone. Develop parameters to characterize this.

D. Establish the system levels at which fault masking can/cannot be tolerated while still allowing the overall system to recognize an upset event.

E. Establish a list of "obvious" digital architecture design practices which can be followed to make a system more upset-tolerant. Notice of the tradeoffs of using these practices vs. other (possibly more classical design techniques e.g. straightforward state minimization) will be given.

F. Formulate a method whereby one can characterize the capacity to failures and environmental pheonomenon present during a lightning event.

G. Development of a fault injection strategy based on experimental data whereby faults can be injected into either real devices or into emulation/simulation of real devices.

H. Refine and extend theory for comparing candidate designs.

REFERENCES

A) J. E. Nanevicz, E. F. Vance, "Analysis of Electrical Transients Created by Lightning," Final Report, Contract NAS1-13792, SRI Project 4026, SRI International, Park, California 94025.

B) E. F. Vance, "Electromagnetic Interference Control," IEEE, EMC-22, No. 4, November 1980.

C) E. F. Vance, "EMP Hardening of Systems," presented at the Fourth EMC Symposium, Zurich, 1981.